SPARSE SUPERPIXEL UNMIXING FOR EXPLORATORY ANALYSIS OF CRISM HYPERSPECTRAL IMAGES

David R. Thompson, Rebecca Castaño *

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Dr.,
Pasadena, CA 91109

Martha S. Gilmore, †

Wesleyan University
Wesleyan Station
Middletown, Connecticut 06459

ABSTRACT

Fast automated analysis of hyperspectral imagery can inform observation planning and tactical decisions during planetary exploration. Products such as mineralogical maps can focus analysts' attention on areas of interest and assist data mining in large hyperspectral catalogs. In this work, sparse spectral unmixing drafts mineral abundance maps with Compact Reconnaissance Imaging Spectrometer (CRISM) images from the Mars Reconnaissance Orbiter. We demonstrate a novel "superpixel" segmentation strategy enabling efficient unmixing in an interactive session. Tests correlate automatic unmixing results based on redundant spectral libraries against hand-tuned summary products currently in use by CRISM researchers.

Index Terms— Sparse Bayesian Unmixing, CRISM, Hyperspectral Images, Superpixels, Image Segmentation

1. INTRODUCTION

Planetary exploration scenarios hold special challenges for hyperspectral image analysis. With the exception of isolated landing sites the ground-truth surface composition is unknown. The number and character of spectral sources is generally uncertain. A trained analyst can guess probable constituents by inspecting individual spectra, but increasing data volumes will preclude comprehensive analysis of this kind. For example, the Compact Reconnaissance Imaging Spectrometer (CRISM) aboard the Mars Reconnaissance Orbiter [1] is collecting images at unprecedented spatial and spectral resolution. It will return over a terabyte of data to Earth over its nominal mission, which is far larger than our capacity for exhaustive manual study.

Timely analysis will require techniques that automatically search an image and summarize possible constituents. With some oversimplification we can categorize current search strategies as "supervised" or "unsupervised." Supervised methods use a detection function, developed on previous data by hand or statistical techniques, to identify one or more specific target signals. For example, hand-selected band ratios can reveal spectral features that are diagnostic of particular minerals [2, 3]. Other supervised detection strategies exploit classifiers such as neural networks or decision trees. These are suited to focused searches for specific mineral types; they may not notice unanticipated or anomalous mineralogy. The detection decision does not always generalize to new scenes and in general only detects constituents that have already been observed and for which training data is available. Alternatively, unsupervised techniques like PCA summarize structure inherent in the spectral data of each new image. Purely unsupervised results are not always physically meaningful, and may ignore subtle or localized mineralogy since they minimize reconstruction error over the entire scene.

Here we examine a spectral unmixing approach to mineralogical search. Our method leverages a redundant library of source spectra from laboratory experiments or previous remote observations. A sparse Bayesian unmixing algorithm computes statistically likely combinations of constituents from this set of candidate sources. A "superpixel" segmentation improves analysis time by orders of magnitude for use in an interactive user session. Preliminary tests suggest that Bayesian spectral unmixing can assist analysis of CRISM data and hyperspectral imagery in general.

2. APPROACH

We use a linear mixing model where m source spectra in w wavelengths combine to yield n observations:

$$\mathbf{x}_\lambda \quad = \quad \mathbf{A}\mathbf{s}_\lambda + \mathcal{N}(0, \sigma^2) \tag{1}$$

Each wavelength λ is associated with an $n \times 1$ observation vector $\mathbf{x}_\lambda$, and an $m \times 1$ source vector $\mathbf{s}_\lambda$. Here $\mathbf{A}$ is an $n \times m$ mixing matrix whose entries describe the contribution of each source signal to the resulting observation. We assume additive zero-mean Gaussian noise of variance σ^2. More generally one can treat the entire spectrum of independent wavelengths as columns of an $m \times w$ source matrix $\mathbf{S}$ and an $n \times w$ observation matrix $\mathbf{X}$. We will index matrix entries with subscripts

* *david.r.thompson@jpl.nasa.gov, rebecca.castano@jpl.nasa.gov*
† *mgilmore@wesleyan.edu*

so that the ith observation on bandwidth λ is $\mathbf{X}_{i\lambda}$. Given source and observation matrices, the unmixing problem aims to recover the mixing matrix $\mathbf{A}$.

2.1. Sparse Bayesian Unmixing

Our geographic linear mixing model places several constraints on $\mathbf{A}$. Minerals' reflectances represent fractions of incident illumination and combine in proportion to their abundance on the imaged surface. To reflect this physical intuition, entries of A must be zero or positive. Standard non-negative unmixing strategies like Nonnegative Matrix Factorization minimize squared reconstruction error while enforcing positivity [4]. In our case there is another important consideration. The source library may contain many more possible sources than actually exist in the scene, so we favor *sparse* solutions where most entries of $\mathbf{A}$ are zero.

We enforce both non-negativity and sparsity using an exponential prior distribution $\mathcal{G}$ parameterized with hyperparameter α, on entries of the mixing matrix. From Bayes' rule:

$$p(\mathbf{A}|\mathbf{X},\alpha) \propto p(\mathbf{X}|\mathbf{A})p(\mathbf{A}|\alpha) \tag{2}$$

$$\propto \prod_{i\lambda} \mathcal{N}(\mathbf{X}_{i\lambda} - [A\mathbf{S}]_{i\lambda}, \sigma^2) \prod_{ij} \mathcal{G}(\mathbf{A}_{ij}, \alpha) \tag{3}$$

$$\propto \prod_{i\lambda} \exp\left\{\frac{-(\mathbf{X}_{i\lambda} - [A\mathbf{S}]_{i\lambda})^2}{2\sigma^2}\right\} \prod_{ij} \exp\{-\alpha\mathbf{A}_{ij}\} \tag{4}$$

We maximize this probability with an iterative gradient descent procedure similar to the Bayesian Positive Source Separation approach described by Moussaoui *et al.* [5]. Taking $\log p(\mathbf{A}|\mathbf{x},\alpha)$ and dropping constant terms produces the objective function $\mathcal{V}(\mathbf{A})$:

$$\mathcal{V}(\mathbf{A}) \quad = \quad -\frac{1}{2\sigma^2}\sum_{i\lambda}(\mathbf{X}_{i\lambda} - [A\mathbf{S}]_{i\lambda})^2 - \sum_{ij}\alpha\mathbf{A}_{ij} \tag{5}$$

Note that this amounts to a least-squares error minimization with an L1-norm penalty term, similar to the sparsity-inducing LASSO estimator [6]. We maximize this with iterative ascent of the gradient of $\mathcal{V}(\mathbf{A})$ with respect to $\mathbf{A}$:

$$\nabla_{\mathbf{A}_{ij}}\mathcal{V}(\mathbf{A}) = [(\mathbf{X} - \mathbf{A}\mathbf{S})\mathbf{S}^T]_{ij} - \alpha \tag{6}$$

The density of the exponential distribution is zero for negative values of $\mathbf{A}$. Our subgradient-inspired approach replaces negative gradient values with zero for entries $\mathbf{A}_{ij}$ within numerical tolerance of zero. This produces an $n \times m$ matrix $\Psi(\mathbf{A})$ representing the update direction:

$$\Psi_{ij}(\mathbf{A}) \quad = \quad \begin{cases} \nabla_{\mathbf{A}_{ij}}\mathcal{V}(\mathbf{A}) & \text{if} \mathbf{A}_{ij} \geq 0 \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

The iterative update at time t follows the update direction Φ^t for the positive step size β:

$$\mathbf{A}^t = \mathbf{A}^{t-1} + \beta\Phi^t \tag{8}$$

At each iteration a line search identifies the optimal step size β. Finally we estimate the noise parameter σ from the reprojected observations as in Moussaoui *et al* [5], fitting noise and mixing parameters in turn until reaching a local optimum. This strategy generally identifies a local Maximum A Priori estimate in less than 100 iterations, permitting real-time spectral unmixing of user-selected image regions.

2.2. Image Segmentation

Segmentation is a valuable tool for mineralogical search. Our goal is to draft compositional maps for large scale data mining of trends and detection of novel mineralogy. This requires analyzing entire images containing thousands or millions of distinct spectra, and spectral unmixing algorithm is still too slow to exhaustively unmix all pixels during an interactive session. Segmentation permits a single mean spectrum to stand in for many pixels, potentially improving run time by orders of magnitude. In addition, segmentation counters pixel-level noise that might otherwise produce false positive detections. Real (non-noise) signals generally come from objects such as outcrops that subtend several adjacent pixels in the high resolution imagery. We can exploit this fact by analyzing the mean spectra of physically-connected regions within the image.

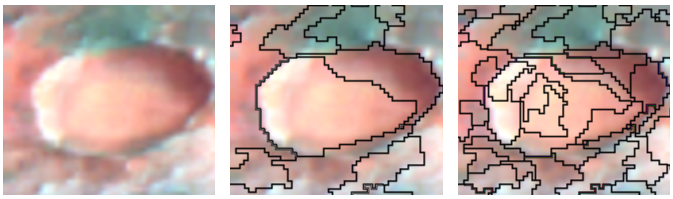

Fig. 1. Segmentation from image 3e12. Finer segmentations provide additional resolution at the cost of greater computation time and sensitivity to noise. Left: Original subimage. Center: coarse segmentation, minimum region size 100. Right: fine segmentation, minimum region size 20.

For mineralogical search it is important that small outlier units of surface material have independent segments. On the other hand one can safely split larger units of surface material into multiple segments. This leads us to intentionally oversegment the scene, a technique known in the Computer Vision community as *superpixel* segmentation. Each superpixel provides some small connected image region that is compositionally homogeneous [7]. Segmentations that produce 3000-5000 superpixels require approximately 20 minutes per image with a modern desktop processor.

We compute superpixels with a graph segmentation. The pixel grid defines an 8-connected graph of vertices and edges. We define edge weights $d(\mathbf{X}_i, \mathbf{X}_k)$ as the sum of squared differences in all wavelengths between neighboring pixels:

$$d(\mathbf{X}_i, \mathbf{X}_k) = \sum_{\lambda}(\mathbf{X}_{i\lambda} - \mathbf{X}_{k\lambda})^2 \tag{9}$$

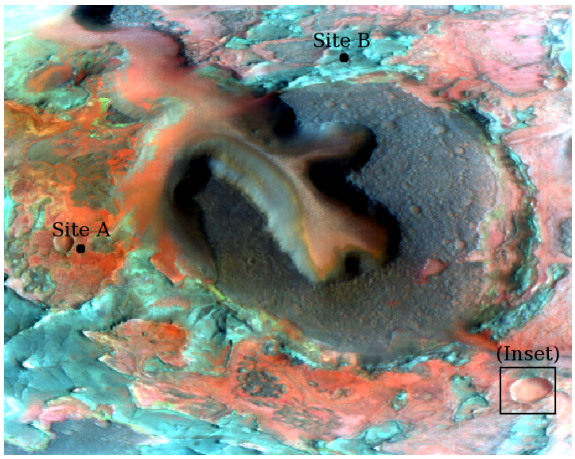

Fig. 2. CRISM 3e12 ($R : 2.0\mu m, G : 1.5\mu m, B : 1.1\mu m$).

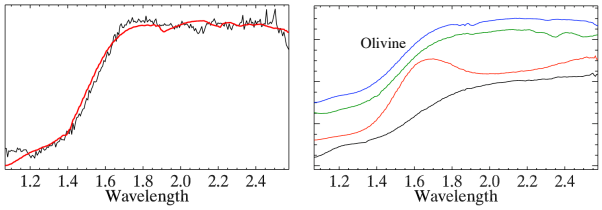

Fig. 3. Site A reconstruction (Left) and constituents (Right).

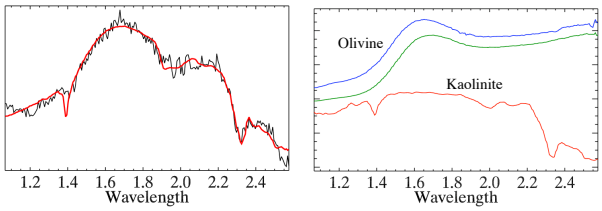

Fig. 4. Site B reconstruction (Left) and constituents (Right).

We cluster pixels by merging them into successively larger subgraphs using the technique of Felzenszwalb *et al* [8]. The Felzenszwalb criterion merges neighboring subgraphs whenever the weight associated with their smallest connecting edge is some threshold larger than the minimum *internal* weight in either subgraph. A final postprocessing step merges regions smaller than a minimum size. Important advantages of this algorithm are efficiency and the ability to segment hyperspectral data based on a hyperspectral distance measure. It can trade speed for accuracy by changing merging thresholds and minimum region size (Figure 1).

3. EVALUATION

A case study analyzes spectra for CRISM Infrared wavelengths from $1.0\mu m$ to $2.5\mu m$. We quantify the correlations for images FRT00003e12 and FRT00003fb9 from the Nili Fossae region (later we omit "FRT0000" prefixes for clarity). Spectra from the first image suggest Olivine and Phyllosilicate minerals [9], while the latter also evidences a strong Olivine signature with evidence for Carbonates such as Magnesite [10]. We preprocess each image using radiometric and atmospheric correction with the Crism Analysis Tool [11].

The source matrix for image 3e12 consists of spectra drawn from the MRO/CRISM spectral library. These include 27 examples of Olivine and 59 examples of phyllosilicates such as Montmorillonite, Nontronite, Saponite, Kaolinite, and Illite. We also augment image 3fb9's library with examples of Magnesite and Hydromagnesite. This diversity helps to account for variation within mineral species. We $L1$-normalize all source spectra during unmixing so that a common hyperparameter α affects all sources equally. Finally, we append 10 featureless line spectra to the library to compensate for arbitrary additive offsets and constant slopes.

Figure 2 shows a typical result for CRISM image 3e12. It identifies sites A and B that show strong signals of olivine and phyllosilicate respectively. Figure 3 (Left) shows a single pixel from site A; the black line is the original spectrum and the red smooth line the reconstruction produced by the sparse mixture. The plot on the right shows the top mixture components. Unsurprisingly each is an instance of olivine from the source library. Figure 4 shows site B with phyllosilicate (Kaolinite, in red) in addition to olivine.

After unmixing one can quantify a mineral's abundance by summing all the mixing coefficients corresponding to its library examples. We use these values to produce a "relative abundance map" that indicates the relative concentrations of the mineral at different superpixel locations. Our evaluation compares these maps to two hand-tuned summary products currently in use by Mars scientists, the OLINDEX and D2300 indicators [2]. These are functions based on slope and band depth that respond strongly to olivine and phyllosilicates respectively. Figures 5 and 6 show relative abundances and the corresponding summary products for image 3e12.

We compute abundance maps for both small and large superpixels (with minimum sizes of 20 and 100 pixels). Table 1 shows correlations between the automated abundance maps and the summary products for the two segmentations. The columns, from left to right, show: the mineral type and index used, the CRISM image, the number of superpixels for coarse (C) and fine (F) segmentations, the linear correlation between abundance maps and the selected index, the Spearman's ρ rank correlation coefficient, the precision score, and the recall score. These last two scores are produced by thresholding summary indices and abundance maps at an appropriate level to yield a binary detection decision at each pixel. Precision considers the fraction of automatic detections that are actually present in the summary product, while recall describes the percentage of summary product detections that are automatically recovered. We advise caution in interpreting

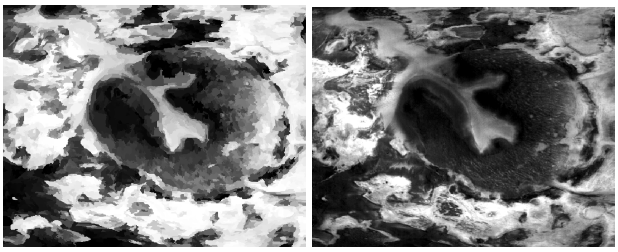

Fig. 5. Olivine abundance (Left) and OLINDEX (Right)

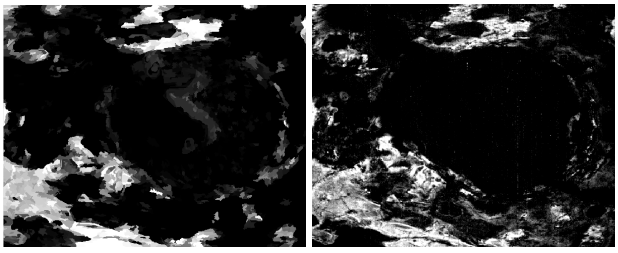

Fig. 6. Phyllosilicate abundance (Left) and D2300 (Right)

Index	Image	n	corr	ρ	prec	rec
Olivine	3e12	(C) 664	0.87	0.91	0.89	0.81
(OLIND)		(F) 3667	0.90	0.95	0.92	0.83
	3fb9	(C) 594	0.87	0.90	0.91	0.86
		(F) 3676	0.92	0.94	0.94	0.87
Phyllosil.	3e12	(C) 664	0.67	0.46	0.76	0.55
(D2300)		(F) 3667	0.73	0.49	0.80	0.53

Table 1. Comparative evaluation (see text for details).

these values since the summary products do not constitute a ground-truth standard to determine mineral abundances and in the case of discrepancies either method could be erroneous.

4. DISCUSSION

The unmixing strategy achieves high correlation with the specific summary products despite being an automated, general method. The abundance maps are fully interpretable through the linear mixing model and component minerals. Lower correlation and Spearman ρ scores for the Phyllosilicate indicator reflect the smaller percentage of this material in the image; irrelevant low-abundance superpixels dominate the score. Umixing successfully detects carbonate in image 3fb9, but we forgo a comparison since standard carbonate summary products use features outside our $1\mu m - 2.5\mu m$ range.

We have presented preliminary results from an unmixing approach to mineralogical search and survey in large image catalogs. Exponential priors encourage a sparse unmixing solution and a graph-based "superpixel" segmentation compresses high-resolution hyperspectral images for improved speed. Rather than relying exclusively on particular diagnostic slopes or band depths, which may be at or below the level of noise, the unmixing approach generates mixtures to explain the entire spectral waveshape. This may eventually improve detection sensitivity beyond that offered by static decision rules. Sparse superpixel unmixing is promising for general automated analysis of hyperspectral image catalogs.

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, through an agreement with the National Aeronautics and Space Administration with support from the Advanced Multi-Mission Operations System (AMMOS).

5. REFERENCES

[1] S. Murchie et al., "Compact reconnaissance imaging spectrometer for Mars (CRISM) on Mars reconnaissance orbiter," *J. Geophys. Res*, vol. 112, no. 5, 2007.

[2] S.M. Pelkey et al., "CRISM multispectral summary products: Parameterizing mineral diversity on Mars from reflectance," *J. Geophys. Res*, vol. 112, 2007.

[3] G. W. Patterson et al., "Developing Tools to Highlight the Presence of Carbonates in CRISM Images of Mars," in *Lunar and Planetary Science Conference*, 2009.

[4] D.D. Lee and H.S. Seung, "Algorithms for non-negative matrix factorization," *Advances in neural information processing systems*, pp. 556–562, 2001.

[5] S. Moussaoui et al., "A Bayesian method for positive source separation," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2004.

[6] C. Leng, Y. Lin, and G. Wahba, "A note on the lasso and related procedures in model selection," *Statistica Sinica*, vol. 16, no. 4, pp. 1273, 2006.

[7] X. Ren and J. Malik, "Learning a classification model for segmentation," in *IEEE International Conference on Computer Vision*, 2003, pp. 10–17.

[8] P.F. Felzenszwalb and D.P. Huttenlocher, "Efficient graph-based image segmentation," *International Journal of Computer Vision*, vol. 59:2, pp. 167–181, 2004.

[9] J.F. Mustard et al., "Hydrated silicate minerals on Mars observed by the Mars Reconnaissance Orbiter CRISM instrument," *Nature*, vol. 454, no. 7202, 2008.

[10] B.L. Ehlmann et al., "Orbital Identification of Carbonate-Bearing Rocks on Mars," *Science*, vol. 322, no. 5909, pp. 1828, 2008.

[11] F. Morgan et al., "CAT tutorial," *CRISM Data User's Workshop, Lunar and Planetary Science Conf.*, 2009.